M.E. ELECTRONICS AND TELE-COMMUNICATION ENGINEERING
FIRST YEAR FIRST SEMESTER EXAM 2024
OPERATING SYSTEMS (COMP)

Time: Three Hours | Full Marks: 100

*Answer **any five** question*
All the sub-parts of the questions should be written altogether

1. a) A computer system uses the Banker's Algorithm to deal with deadlocks. Its current state is shown in the tables below, where P0, P1, P2 are processes and R0, R1, R2 are resource types.

	Maximum Need				Current Allocation			Available		
	R0	R1	R2		R0	R1	R2	R0	R1	R2
P0	4	1	2	P0	1	0	2	2	2	0
P1	1	5	1	P1	0	3	1			
P2	1	2	3	P2	1	0	2			

i. Show that the system can be in this state.
ii. What will the system do on a request by process P0 for one unit of resource type R1?

b) Consider the resource allocation graph given in the figure.

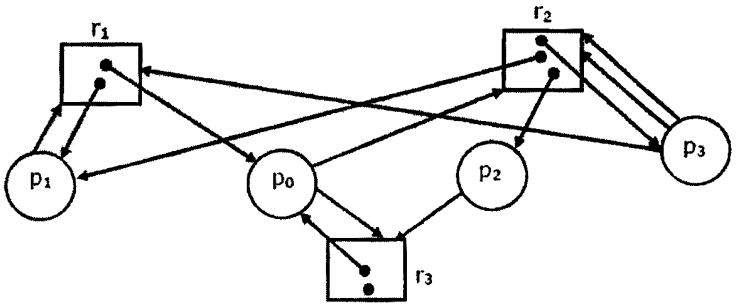

i. Find if the system is in a deadlock state.
ii. Otherwise, find a safe sequence.

c) A single processor system has three resource types X, Y and Z, which are shared by three processes. There are 5 units of each resource type. Consider the following scenario, where the column alloc denotes the number of units of each resource type allocated to each process, and the column request denotes the number of units of each resource type requested by a process in order to complete execution. Which of these processes will finish LAST?

	alloc			request		
	X	Y	Z	X	Y	Z
P0	1	2	1	1	0	3
P1	2	0	1	0	1	2
P2	2	2	1	1	2	0

8+8+4 = 20

2. a) Define various performance criteria used for comparing CPU scheduling algorithms.
b) Assume you have the following processes to execute with one processor. All five process arrive at time 0, in the order given, with the length of the CPU – burst time given in miliseconds.

Process	Burst Time
P1	**10**
P2	**29**
P3	**3**
P4	**7**
P5	**12**

Consider the FCFS, SJF and RR (quantum = 10 miliseconds) scheduling algorithms for this set of processes. Which algorithm would give the minimum average waiting time?

10+10 = 20

3. a) Explain Process Control Block. What is Dynamic Linked Libraries?
b) What do you mean by page-faults? When do page-faults occur? Describe the action taken by the O.S. when page fault occurs.
c) Two processes, P1 and P2, need to access a critical section of code. Consider the following synchronization construct used by the processes: Here, wants1 and wants2 are shared variables, which are initialized to false. Which one of the following statements is TRUE about the above construct?

```
/* P1 */                          /* P2 */
while (true) {                    while (true) {
 wants1 = true;                    wants2 = true;
 while (wants2 == true);           while (wants1==true);
 /* Critical                       /* Critical
   Section */                        Section */
 wants1=false;                     wants2 = false;
}                                 }
/* Remainder section */           /* Remainder section */
```

i. It does not ensure mutual exclusion. ii. It does not ensure bounded waiting.
iii. It requires that processes enter the critical section in strict alternation.
iv. It does not prevent deadlocks, but ensures mutual exclusion. 9+9+2 = 20

4. a) What is CPU scheduling? What is its need? List various scheduling algorithms.
b) What is fragmentation? Explain the difference between internal and external fragmentation.
c) Consider three processes, all arriving at time zero, with total execution time of 10, 20 and 30 units, respectively. Each process spends the first 20% of execution time doing I/O, the next 70% of time doing computation, and the last 10% of time doing I/O again. The operating system uses a shortest remaining compute time first scheduling algorithm and schedules a new process either when the running process gets blocked on I/O or when the running process finishes its compute burst. Assume that all I/O operations can be overlapped as much as possible. For what percentage of time does the CPU remain idle?

8+8+4 = 20

5. a) State and discuss the differences between Multiprogramming and Multitasking.
   b) Given references to the following pages by a program:

   0,9,0,1,8,1,8,7,8,7,1,2,8,2,7,8,2,3,8,3

   How many page faults will occur if the program has three frames available to it and uses
   - i. First-in-First Out (FIFO) replacement strategy.
   - ii. Least-Recently-Used (LRU) replacement strategy.
   - iii. Optimal Replacement strategy.

   c) Consider a paging system that uses 1-level page table residing in main memory and a TLB for address translation. Each main memory access takes 100 ns and TLB lookup takes 20 ns. Each page transfer to/from the disk takes 5000 ns. Assume that the TLB hit ratio is 95%, page fault rate is 10%. Assume that for 20% of the total page faults, a dirty page has to be written back to disk before the required page is read from disk. TLB update time is negligible. The average memory access time in ns (round off to 1 decimal places) is _____. 8+8+4 = 20

6. a) Define a critical section and the critical section problem, and explain the requirements of a solution to the critical section problem.
   b) Describe the paging scheme of memory management with reference to the hardware required to support the scheme.
   c) Consider a machine with 64 MB physical memory and a 32-bit virtual address space. If the page size is 4KB, what is the approximate size of the page table? 8+8+4 = 20

7. a) Consider the dining philosophers problem. If we add a 6th chopstick to the center of the table, have we cured the deadlock problem? If yes, what condition have we removed? If no, explain why not.
   b) Suppose the following disk request sequence (track numbers) for a disk with 100 tracks is given: 45, 20, 90, 10, 50, 60, 80, 25, 70. Assume that the initial position of the R/W head is on track 50. The additional distance that will be traversed by the R/W head when the Shortest Seek Time First (SSTF) algorithm is used compared to the SCAN (Elevator) algorithm (assuming that SCAN algorithm moves towards 100 when it starts execution) is _____ tracks. 10+10 = 20

8. Write notes on the following:
   a) Life Cycle of a Process.
   b) Semaphore as a synchronization tool.
   c) Resource Allocation Graph.
   d) Kernel of an Operating System. 4x5 = 20